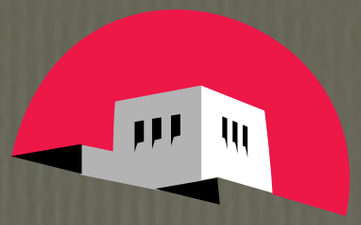

The University of New Mexico

Complexity Issues

Barney Maccabe
Computer Science Department
&
Center for High Performance Computing

HPC@UNM

The Center for High Performance Computing

Bandwidth

- In the next decade is the bandwidth transferred into or out of one “high end computing file system”
  a. going down 10X or more,
  b. staying about the same,
  c. going up 10X or more, or
  d. “your answer here,”

  as a result of the expected increase in computational speed in its client clusters/ MPPs, and why?

Increased Compute Rate...

- .... is **not** the only reason for the increased bandwidth requirement

- External sources/sinks will also create and consume data at a higher rate
- Growth rate of aggregate MPP memory

Aggregate Memory Growth

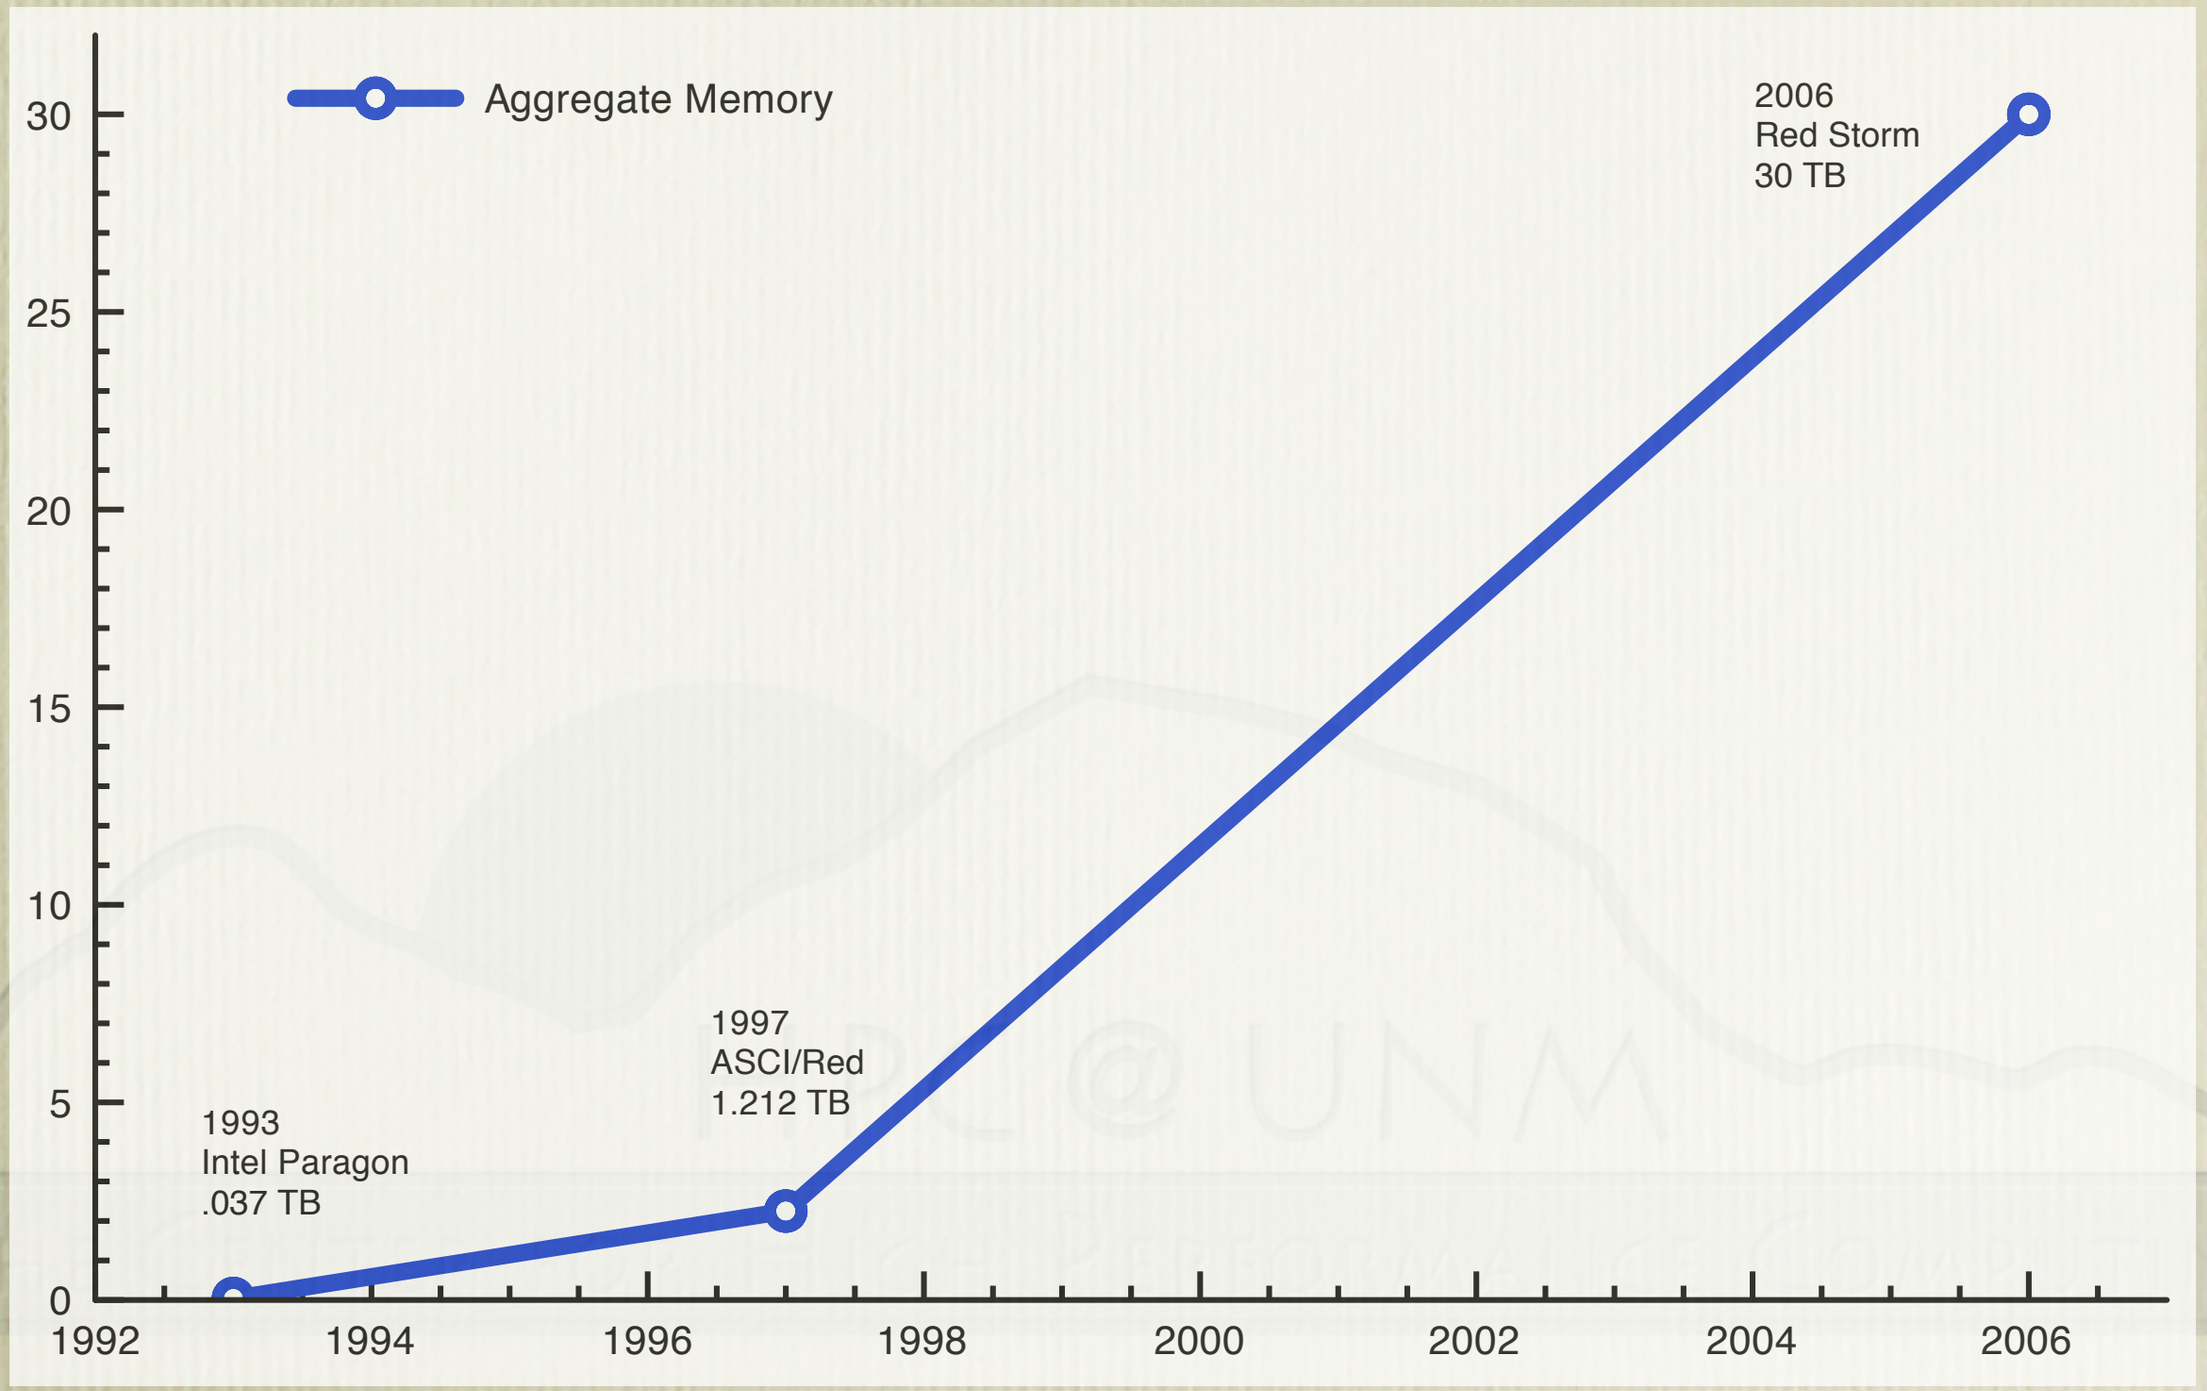

Spindle Count

- In the next decade is the number of magnetic disks in one “high end computing file system”
    c. going up 10X or more, or
  as a result of the expected increase in computational speed in its client clusters/ MPPs, and why?

- It’s unlikely that increases in disk density will be able to keep up with the increased demand for storage

Others

- Concurrency: the number of concurrent streams of requests is likely to increase by at least 10x
  - more systems, more users, more remote access, ...
- Seek Efficiency: I have no clue how the number of bytes moved per magnetic disk seek will change
- Failures: the number of independent failure domains **may** increase by 10x
  - we don't tolerate failures very well.

Complexity

- Explain why these large increases are not going to increase the complexity of storage software significantly.
- Make the user deal with it

Development Time Trends

The University of New Mexico

- Even if complexity is increasing, the time and effort required to achieve acceptable 9s of availability at speed **MUST** stay about the same.
- Are you relying on the development of any currently insufficient technologies, and if so, which?
  - **magic**?

The temptation

- Fat layers provide opportunities for optimization
  - they also provide lots of opportunities to do the wrong thing

- Fat layers are magic–magic is (mostly) OK on my laptop...

- Getting to fat layers requires a lot of experience

The obligatory F1 analogy

The University of New Mexico

Shifting gears

Falcon	Synchromesh	Automatic Transmission
F1	Heal & toe	Automatic Clutch
	skill	*enhancement*

- Wired Magazine; March 2001
  - 220 MPH, 17,000 RPM, 500,000 lines of code

Performance Transparency

- When performance matters, the API should accurately reflect resource costs
  - FORTRAN reflects costs of the (an?) underlying architecture
  - MPI reflects costs of memory distribution
  - It's hard to encourage people to do the right thing–MPI does not encourage people to overlap communication with computation
- Let users see through and discard layers

Coping with complexity

The University of New Mexico

Butler Lampson (Software 1, 1), simplicity:

Perfection is reached, not when there is no longer anything to add, but when there is no longer anything to take away.

A. Saint-Exupery

- Make it fast, rather than general or powerful
- Don't hide power
- Leave it to the client

Abstracting storage

- You can't fix latency, you have to tolerate it
  - asynchronous interface (split transactions)
  - find something else to do
- I/O Bandwidth needs to be manageable
  - expose parallelism
  - eliminate need for anything other than data movement and **access control**
  - provide mechanisms to do everything else "offline"
- Mask failures

LWFS

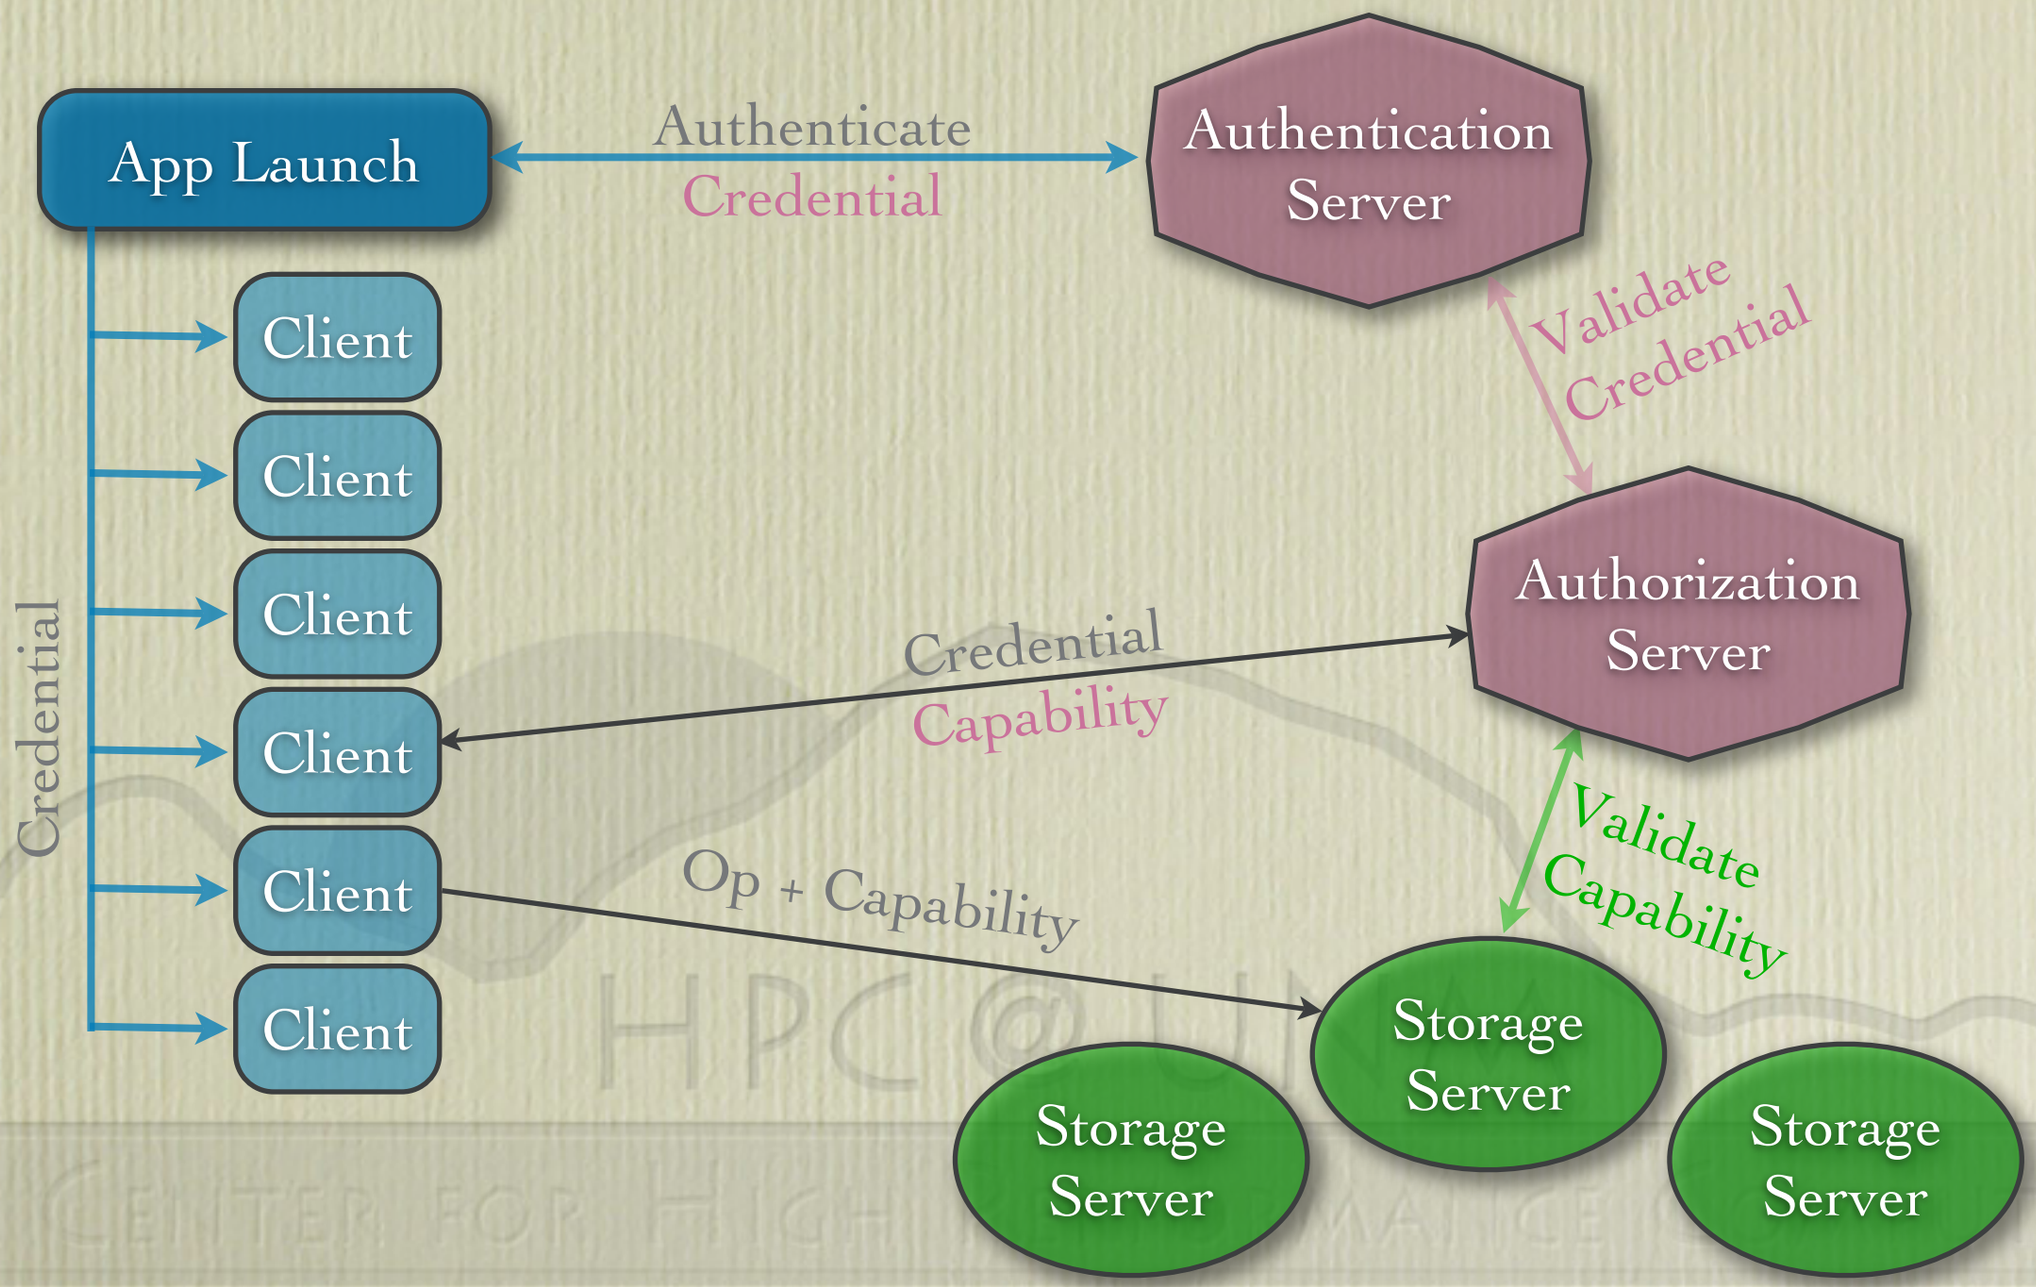
App Launch
Authenticate
Credential
Authentication Server
Client
Client
Client
Client
Client
Client
Credential
Validate Credential
Authorization Server
Credential
Capability
Validate Capability
Op + Capability
Storage Server
Storage Server
Storage Server

Why Data isn't Compute

- Persistence
  - rebooting the storage system won't clean it up
- Sharing
  - performance isolation is much harder–you can't space share the storage servers
- Interdependencies
  - management requires the construction of global views